Quantifying Human Priors over Social and Navigation Networks

Gecia Bravo-Hermsdorff

Previous: Princeton Neuroscience Institute, Princeton University *Current:* Department of Statistics, University College London

TL;DR: To quantify human priors over the structure of connections, we overcame three main hurdles

Engaging human attention.
For experiments that essentially ask participants to be creative, it is important to make the experiment engaging and streamlined.

Parameterizing the priors.
For “smooth” distributions over the space of graphs, a good choice is to maximize entropy given the density of subgraphs with $\leq r$ edges as constraints.

Summarizing the distributions.
To characterize and compare distributions over graphs, graph cumulants are very intuitive.

Brains rely on efficient priors

Neural representations have adapted to efficiently encode the relevant statistics of our environment.

Studying the corresponding priors has lead to foundational results in neuroscience and artificial intelligence.

Just as visual priors “color” our perception...

(Graham & Field, 2009)

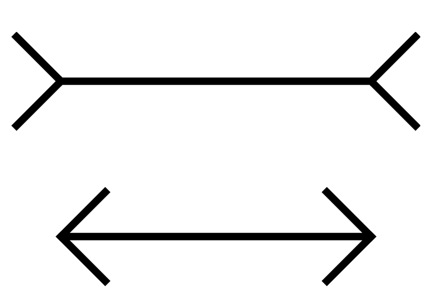

(Howe & Purves, 2004)

...might there be similar “illusions” that exploit our priors over connections?

We focus on social and navigation networks

Two domains that have been quotidian over evolutionary timescales...

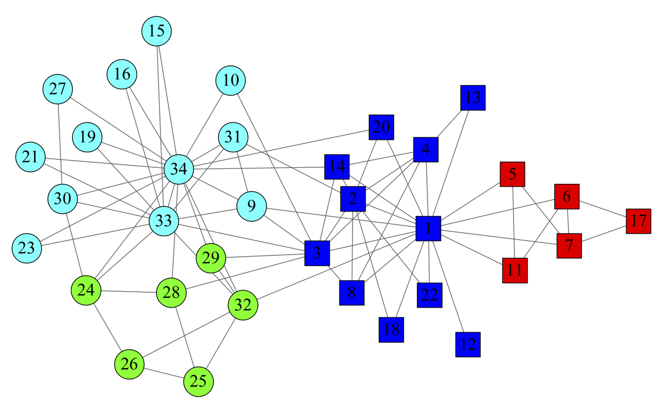

...and are naturally represented as graphs.

Distinguishing between social and navigation

When the model uses subgraphs with more than 3-4 edges...

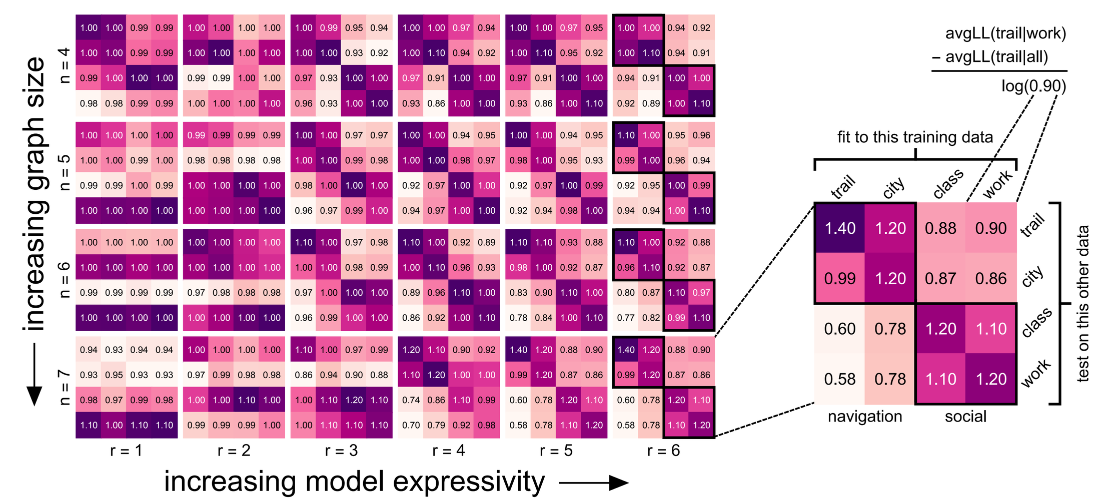

...generalization within a domain is notably better than across different domains.

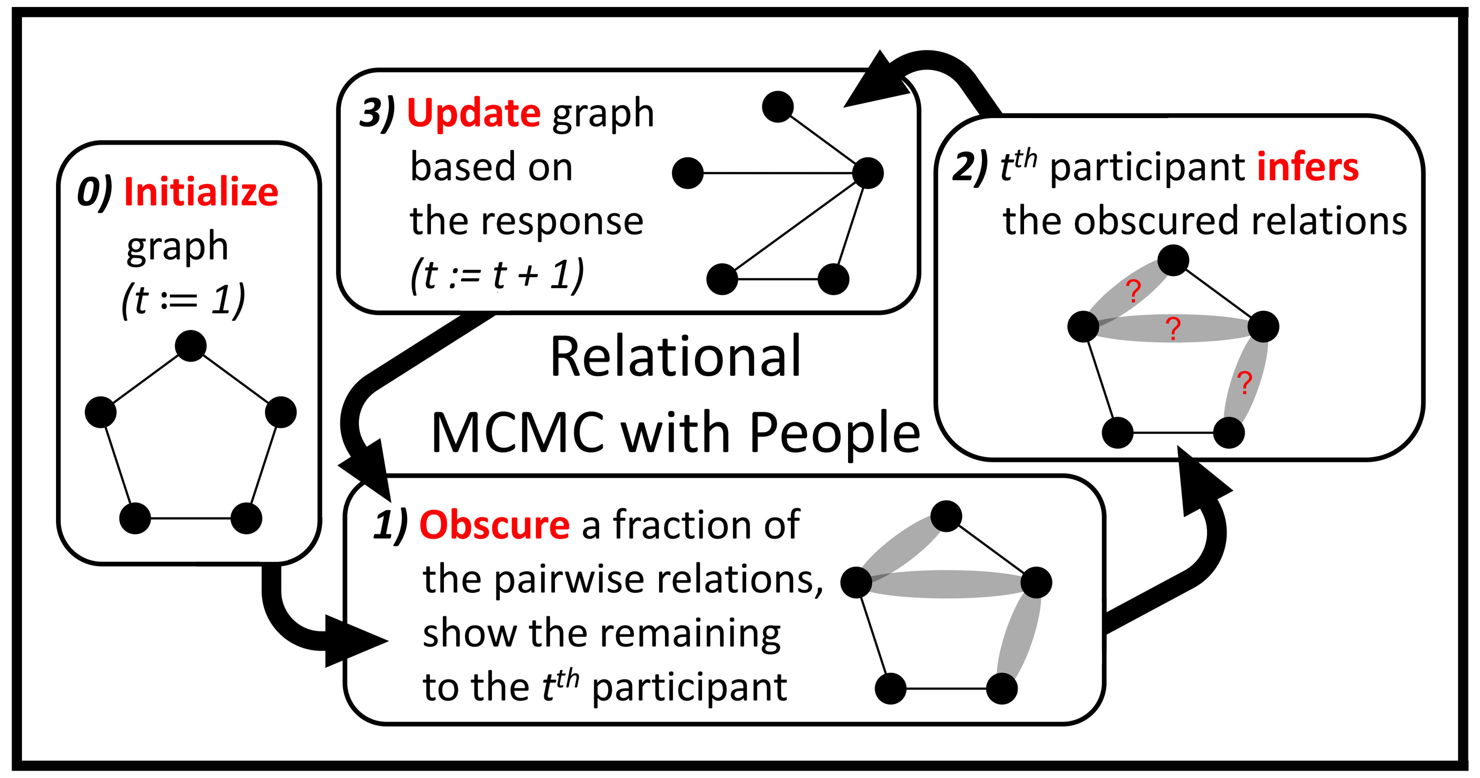

An intuitive interface and engaging cover stories

domain	context	nodes	relations
social	class	students	friendships
	work	coworkers	friendships
navigation	city	neighborhoods	borders
	park	nature sites	trails

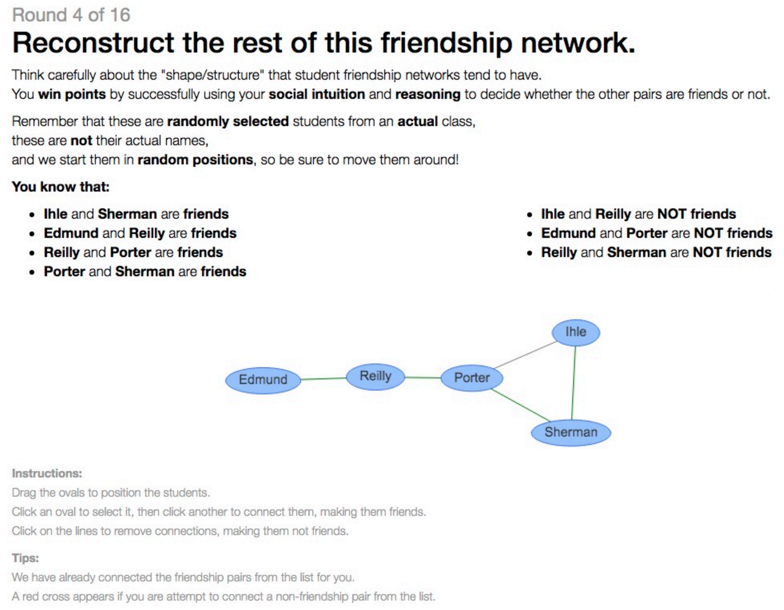

We found it was helpful to have: compelling cover stories, carefully worded instructions, and an interactive interface.

Sparsity, heterogeneity, and clustering

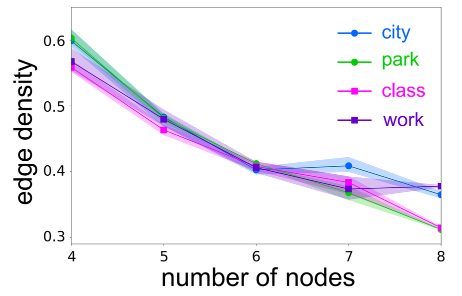

Priors over larger graphs have lower edge density.

Priors over smaller graphs have fewer hubs.

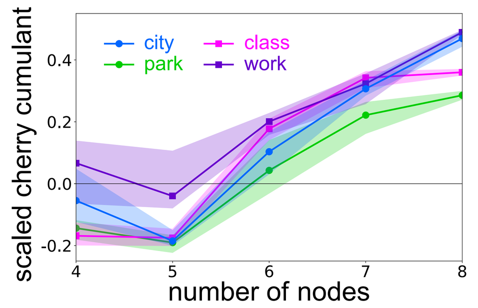

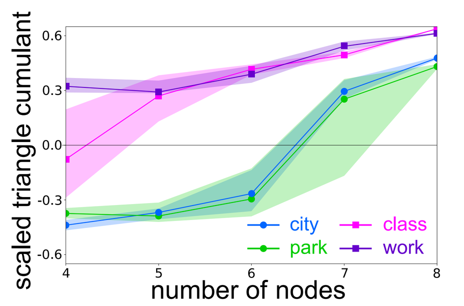

Priors over social graphs have more triangles.

Graph cumulants are intuitive summary statistics

Mathematically analogous to the classical cumulants (mean, variance, skew, kurtosis, etc)...

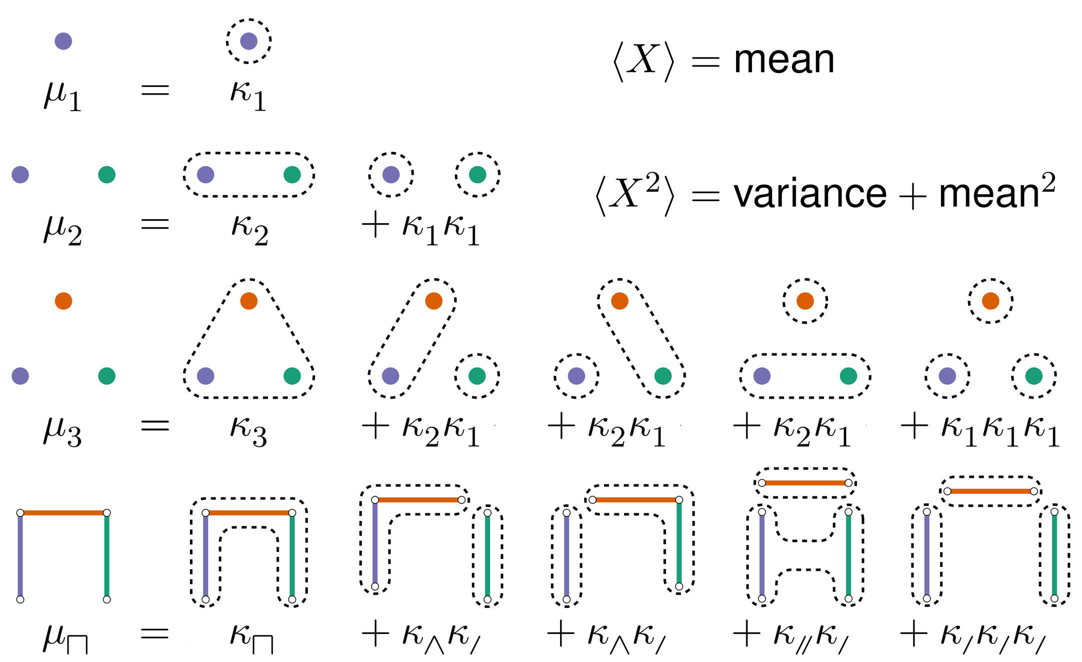

$\langle X \rangle$ = mean

$\langle X^2 \rangle$ = variance + mean^2

...graph cumulants quantify “excess propensity” for a subgraph.

Maximum-entropy distributions for modelling priors

Given a set of subgraphs with prescribed subgraph densities $\vec{\mu}_g$, the maximum-entropy distribution is of the form:

$$\pi(G) = \mathrm{ER}_{n,1/2}(G) \times \exp \sum_g \beta_g \mu_g(G)$$

A Bayesian agent with prior π, presented with a “partial graph” PG_t will sample a graph G_t from their posterior distribution:

$$p(G_t|PG_t) = \frac{p(PG_t|G_t)\pi(G_t)}{\sum_G p(PG_t|G)\pi(G)}$$

To find $\vec{\beta}_g$, we compute the log-likelihood of the data...

$$\mathcal{L}(\vec{\beta}) = \sum_t \log \frac{\mathrm{ER}_{n,1/2}\Big(G_t\Big|PG_t\Big)\ \exp\Big\{\vec{\beta}\cdot\vec{\mu}(G_t)\Big\}}{\sum_{G'\in\mathcal{G}_n} \mathrm{ER}_{n,1/2}\Big(G'\Big|PG_t\Big)\ \exp\Big\{\vec{\beta}\cdot\vec{\mu}(G')\Big\}}$$

...and apply Newton’s method until convergence.